

Quantum computational gradient estimation

David W. Bulger*

February 1, 2008

Abstract

Classically, determining the gradient of a black-box function $f : \mathbb{R}^p \rightarrow \mathbb{R}$ requires $p + 1$ evaluations. Using the quantum Fourier transform, two evaluations suffice. This is based on the approximate local periodicity of $e^{2\pi i \lambda f(x)}$. It is shown that sufficiently precise machine arithmetic results in gradient estimates of any required accuracy.

MSC2000 Subject Classification: 90C30, 68Q99, 68Q25

Key words and phrases: quantum computation, gradient estimation.

1 Introduction

The vector gradient of a real-valued function f of a vector argument can be calculated using just two calls to a black-box quantum oracle for f . The mechanism is simple, and capitalises on the fact that, in the vicinity of a point x , $e^{2\pi i \lambda f(x)}$ is periodic, with period parallel and inversely proportional to $\nabla f(x)$. A superposed state is created discretising a small hyperrectangle around the domain point, the function is evaluated, the phase is rotated in proportion to the function value, the oracle call is reversed, and a multidimensional quantum Fourier transform is applied to the bits encoding the discretised hyperrectangle.

This paper establishes, under mild conditions on f , that the gradient estimation can be performed to any required level of accuracy, in the sense that, given any $\delta > 0$ and $\epsilon < 1$, we can produce a superposition of gradient estimates which, if observed, will collapse to an estimate within δ of the true gradient with probability at least ϵ . Greater accuracy is achieved by increasing arithmetic precision and by increasing the number of points in the sampling grid.

The paper's structure is as follows. Section 2 presents some assumptions on the function f whose gradient is sought, and Section 3 formalises the evaluation and manipulation of values of f within the quantum computer. Section 4 presents the gradient estimation algorithm. Section 5 analyses the effect of the algorithm, consisting mostly of the statement and proof of the main result, that any required accuracy is attainable by using sufficiently precise arithmetic and a large enough sampling grid. Because the rest of paper discusses the computations rather abstractly, Section 6 briefly comments on how the algorithm would be performed in practice.

*Lecturer, Department of Statistics, Macquarie University, NSW 2109, Australia.
dbulger@efs.mq.edu.au.

2 Problem formulation

Let $D \subset \mathbb{R}^p$ have non-empty interior. Let f be a twice-differentiable function from D to \mathbb{R} . At each point $x \in D$, let $\nabla f(x)$ denote the gradient and $Hf(x)$ the Hessian matrix of f . Assume that $\|\nabla f(x)\|_\infty \leq L$ and $\|Hf(x)\|_2 \leq M$ for all $x \in D$. It is desired to determine $\nabla f(x)$ for a point x in the interior of D .

3 Oracle Formalism

This paper's main result is that an objective function's gradient can be calculated to any desired precision using the quantum algorithm described. Clearly, any particular data encoding method will only support a certain maximum precision; we therefore require a formalism in which points in the domain and range of f can be represented in a variety of ways. This section introduces the Hilbert spaces and operators involved.

It will be helpful firstly to catalog the operators to be used as they would look if precision and rounding error were *not* relevant. The computational system is a tripartite system; the three parts have state spaces \mathcal{D} , \mathcal{R} and \mathcal{G} (standing for 'domain', 'range' and 'grid'), so the combined system has state space $\mathcal{D} \otimes \mathcal{R} \otimes \mathcal{G}$. Each computational basis state is a tensor product $|d\rangle \otimes |r\rangle \otimes |g\rangle$ of one computational basis state from each of the three factor spaces. For now, suppose that the basis indices d , r and g belong respectively to D , \mathbb{R} and \mathbb{R}^p . Let g_0 be the constant length- p vector $(2^{n-1} - 1/2, \dots, 2^{n-1} - 1/2)$. The operators are

$$\begin{aligned} U_f(|d\rangle \otimes |r\rangle \otimes |g\rangle) &= |d\rangle \otimes |r + f(d)\rangle \otimes |g\rangle, \\ U_+(|d\rangle \otimes |r\rangle \otimes |g\rangle) &= |d + \mu(g - g_0)\rangle \otimes |r\rangle \otimes |g\rangle, \\ U_R(|d\rangle \otimes |r\rangle \otimes |g\rangle) &= e^{2\pi i \lambda f(d)} |d\rangle \otimes |r\rangle \otimes |g\rangle, \end{aligned}$$

where μ and λ are algorithm parameters, as well as the quantum Fourier transform and inverses of U_f and U_+ .

Return now to precision considerations. Suppose that, for any positive ν and μ and natural n , we can construct a quantum oracle evaluating $f(x + \mu(g - g_0))$ for $g \in \{0, \dots, 2^n - 1\}^p$ and $x, x + \mu(g - g_0) \in D$, with an error uniformly bounded by ν . In particular, suppose that we can construct a system $(U_f, U_+, \mathcal{D}, B_{\mathcal{D}}, \mathcal{R}, B_{\mathcal{R}}, \mathcal{G}, c_d, c_r, c_f, c_p)$, where

- \mathcal{D} , \mathcal{R} and \mathcal{G} are finite-dimensional Hilbert spaces,
- $B_{\mathcal{D}}$ and $B_{\mathcal{R}}$ are orthonormal bases for \mathcal{D} and \mathcal{R} ,
- $B_{\mathcal{G}} = \{0, \dots, 2^n - 1\}^p$ is an orthonormal basis for \mathcal{G} ,
- $B_{\mathcal{R}}$ is a group, under an operation we denote as '+', with an identity element we denote as $|0\rangle$,
- $c_d : D \rightarrow B_{\mathcal{D}}$ (the 'domain encoding function'), $c_f : B_{\mathcal{D}} \rightarrow B_{\mathcal{R}}$, $c_r : B_{\mathcal{R}} \rightarrow \mathbb{R}$ (the 'range decoding function'), and $c_p : B_{\mathcal{D}} \times B_{\mathcal{G}} \rightarrow B_{\mathcal{D}}$,
- U_f and U_+ are unitary operators on $\mathcal{D} \otimes \mathcal{R} \otimes \mathcal{G}$, given by

$$\begin{aligned} U_f|d\rangle \otimes |r\rangle \otimes |g\rangle &= |d\rangle \otimes |r + c_f(d)\rangle \otimes |g\rangle, \\ U_+|d\rangle \otimes |r\rangle \otimes |g\rangle &= |c_p(d, g)\rangle \otimes |r\rangle \otimes |g\rangle, \end{aligned}$$

- c_p acts invertibly on $B_{\mathcal{D}}$, so that for each $d \in B_{\mathcal{D}}$ and $g \in B_{\mathcal{G}}$, $c_p^{-1}(c_p(d, g), g) = d$,
- $|f(x) - c_{r \circ c_f \circ c_d}(x)| \leq \nu/2$ for all $x \in D$,
- $|c_{r \circ c_f \circ c_d}(x + \mu(g - g_0)) - c_{r \circ c_f \circ c_p}(c_d(x), g)| \leq \nu/2$ for all $x \in D$ and $g \in B_{\mathcal{G}}$, provided $x + \mu(g - g_0)$ is also in D ,

and, further, that we can implement U_f and U_+ on a quantum computer. (Note that this formalism does not necessarily require the domain points represented by $B_{\mathcal{D}}$ to form a grid; this may be of interest in optimising functions on manifolds.)

4 Algorithm

The algorithm dealt with in this paper estimates the gradient of f at a point x in the interior of its domain. Firstly, using quantum superposition, f is evaluated at every point of a hyperrectangular grid centred around x . The grid is small enough that f is approximately linear across it. Next, the phase of the quantum computational system is rotated in proportion to the value of f at each grid point. Now the phase varies approximately periodically over the grid, and the period determines $\nabla f(x)$. The period is easily determined by the quantum Fourier transform.

Two of the operators involved in the gradient estimation algorithm, U_f and U_+ , were hypothesised in Section 3. Additionally, we will require their inverses U_f^{-1} and U_+^{-1} , a phase rotation operator U_R , and a p -dimensional quantum Fourier transform U_{QFT} .

The operators U_f^{-1} and U_+^{-1} invert the actions of U_f and U_+ , mapping $|d\rangle \otimes |r\rangle \otimes |g\rangle$ to $|d\rangle \otimes |r - c_f(d)\rangle \otimes |g\rangle$ (the subtraction $r - c_f(d)$ is according to the group structure assumed on $B_{\mathcal{R}}$) and $|c_p^{-1}(d, g)\rangle \otimes |r\rangle \otimes |g\rangle$. Note that the function f is *not* being inverted.

The phase rotation operator U_R involves a parameter $\lambda \in \mathbb{R}$, mapping $|d\rangle \otimes |r\rangle \otimes |g\rangle$ to $e^{2\pi i \lambda c_r(r)} |d\rangle \otimes |r\rangle \otimes |g\rangle$. The multidimensional quantum Fourier transform U_{QFT} acts on \mathcal{G} , mapping $|d\rangle \otimes |r\rangle \otimes |g\rangle$ to $2^{-pn/2} \sum_{h \in B_{\mathcal{G}}} e^{2\pi i h \cdot g / 2^n} |d\rangle \otimes |r\rangle \otimes |h\rangle$.

With these operators defined, the gradient estimation algorithm is easily stated. Firstly, the state $|c_d(x)\rangle \otimes |0\rangle \otimes |0\rangle$ is prepared in $\mathcal{D} \otimes \mathcal{R} \otimes \mathcal{G}$, where $x \in D$ is the point at which ∇f is sought. Then, the system $\mathcal{D} \otimes \mathcal{R} \otimes \mathcal{G}$ is subjected to $U_{QFT} \circ U_+^{-1} \circ U_f^{-1} \circ U_R \circ U_f \circ U_+ \circ U_{QFT}$. This results, as we shall see in Section 5, in a state $|c_d(x)\rangle \otimes |0\rangle \otimes |\chi\rangle$, where in general $|\chi\rangle$ is a superposition of computational basis states from $B_{\mathcal{G}}$.

Interpretation of the resulting state $|\chi\rangle$ involves the “gradient decoding function” $c_g : B_{\mathcal{G}} \rightarrow \mathbb{R}^p$, defined by

$$c_g : (g_1, \dots, g_p) \rightarrow (c_{g,1}(g_1), \dots, c_{g,p}(g_p)), \text{ where}$$

$$c_{g,m}(g_m) = \begin{cases} \frac{-g_m}{2^n \lambda \mu}, & g_m \in \{0, \dots, 2^{n-1} - 1\}, \\ \frac{2^n - g_m}{2^n \lambda \mu}, & g_m \in \{2^{n-1}, \dots, 2^n - 1\}. \end{cases}$$

If $|\chi\rangle$ is a basis state $|g\rangle$, it indicates that $\nabla f(x) = c_g(g)$. If, on the other hand, $|\chi\rangle$ is a superposition $\sum_{g \in B_{\mathcal{G}}} \chi_g |g\rangle$, then the gradient estimate is indeterminate, comprising the various discretised values g with the weights $|\chi_g|^2$.

Altogether, in addition to the argument x , the gradient estimation algorithm depends on the four parameters n , ν , λ and μ . Accordingly, the algorithm will be denoted $A(n, \nu, \lambda, \mu; x)$.

5 Behaviour

The state resulting from the algorithm $A(n, \nu, \lambda, \mu; x)$ is

$$\begin{aligned}
& U_{QFT} \circ U_+^{-1} \circ U_f^{-1} \circ U_{R \circ} U_{f \circ} U_{+ \circ} U_{QFT}(|c_d(x)\rangle \otimes |0\rangle \otimes |0\rangle) \\
&= 2^{-pn/2} \sum_{h \in B_G} U_{QFT} \circ U_+^{-1} \circ U_f^{-1} \circ U_{R \circ} U_{f \circ} U_{+ \circ} (|c_d(x)\rangle \otimes |0\rangle \otimes |h\rangle) \\
&= 2^{-pn/2} \sum_{h \in B_G} U_{QFT} \circ U_+^{-1} \circ U_f^{-1} \circ U_{R \circ} U_f(|c_p(c_d(x), h)\rangle \otimes |0\rangle \otimes |h\rangle) \\
&= 2^{-pn/2} \sum_{h \in B_G} U_{QFT} \circ U_+^{-1} \circ U_f^{-1} \circ U_R(|c_p(c_d(x), h)\rangle \otimes |c_{f \circ} c_p(c_d(x), h)\rangle \otimes |h\rangle) \\
&= 2^{-pn/2} \sum_{h \in B_G} e^{2\pi i \lambda c_{r \circ} c_{f \circ} c_p(c_d(x), h)} U_{QFT} \circ U_+^{-1} \circ U_f^{-1} (|c_p(c_d(x), h)\rangle \otimes |c_{f \circ} c_p(c_d(x), h)\rangle \otimes |h\rangle) \\
&= 2^{-pn/2} \sum_{h \in B_G} e^{2\pi i \lambda c_{r \circ} c_{f \circ} c_p(c_d(x), h)} U_{QFT} \circ U_+^{-1} (|c_p(c_d(x), h)\rangle \otimes |0\rangle \otimes |h\rangle) \\
&= U_{QFT} |c_d(x)\rangle \otimes |0\rangle |\psi\rangle \\
&= |c_d(x)\rangle \otimes |0\rangle \otimes |\chi\rangle,
\end{aligned}$$

where

$$|\psi\rangle = 2^{-pn/2} \sum_{h \in B_G} e^{2\pi i \lambda c_{r \circ} c_{f \circ} c_p(c_d(x), h)} |h\rangle \text{ and} \quad (1)$$

$$|\chi\rangle = U_{QFT} |\psi\rangle, \quad (2)$$

and of course U_{QFT} , acting on \mathcal{G} alone, is defined by

$$U_{QFT}(|g\rangle) = 2^{-pn/2} \sum_{h \in B_G} e^{2\pi i h \cdot g / 2^n} |h\rangle.$$

Theorem 1 *For any $\gamma, \delta > 0$ and $\epsilon < 1$, there exist parameters n, ν, λ and μ such that, at every x with $x + [-\gamma, \gamma]^d \subseteq D$, when $|\chi\rangle$ is produced according to $A(n, \nu, \lambda, \mu; x)$,*

$$\|P|\chi\rangle\|_2 \geq \epsilon, \quad (3)$$

where P is the projection

$$\sum \{|h\rangle\langle h| : \|c_g(h) - \nabla f(x)\|_\infty < \delta\}.$$

Proof It will be demonstrated that, if n, ν, λ and μ are chosen satisfying

$$4^{n-1} \pi \lambda M \mu^2 / \sqrt{5} \leq (1 - \epsilon)/3, \quad (4)$$

$$2\pi \lambda \nu \leq (1 - \epsilon)/3, \quad (5)$$

$$2^{n-1} \mu \leq \gamma, \quad (6)$$

$$1/2\lambda\mu \geq L + \delta \text{ and} \quad (7)$$

$$\csc(\pi \lambda \mu \delta) \leq \sqrt{2^n (1 - ((2 + \epsilon)/3)^{2/p})}, \quad (8)$$

then (3) holds. The reader can verify that one choice satisfying (4) to (8) is

$$\begin{aligned} n &= \left\lceil -\log_2 \left(\sin^2(\pi\delta/2(L+\delta)) \left(1 - ((2+\epsilon)/3)^{2/p} \right) \right) \right\rceil, \\ \lambda &= \max \left\{ \frac{2^{n-2}}{\gamma(L+\delta)}, \frac{3 \times 4^{n-2} \pi M}{\sqrt{5}(L+\delta)^2(1-\epsilon)} \right\}, \\ \mu &= 1/2\lambda(L+\delta), \\ \nu &= (1-\epsilon)/6\pi\lambda. \end{aligned}$$

The algorithm contains three sources of error:

- $\nabla f(x)$ will not, in general, be exactly equal to $c_g(g)$ for some $g \in B_G$;
- ∇f will not, in general, be exactly constant throughout the sampling grid;
- calculations are performed to a finite precision, so that $c_{r \circ c_f \circ c_p}(c_d(x), h)$ will not, in general, exactly equal $f(x + \mu(h - g_0))$.

In the oracle's calculation of $f(x + \mu(h - g_0))$, let $\epsilon_D(x, h)$ represent the error due to computational precision and let $\epsilon_N(x, h)$ represent the departure from linearity, so that $c_{r \circ c_f \circ c_p}(c_d(x), h) = f(x + \mu(h - g_0)) + \epsilon_D(x, h) = f(x) + \nabla f(x) \cdot \mu(h - g_0) + \epsilon_N(x, h) + \epsilon_D(x, h)$.

('D' stands for 'discretisation' and 'N' for 'nonlinear'.) By (6), $x + \mu(h - g_0) \in D$. By assumption, $|\epsilon_D(x, h)| \leq \nu$, and by Lagrange's remainder for Taylor's series, we have $|\epsilon_N(x, h)| \leq M\mu^2(h - g_0) \cdot (h - g_0)/2$.

We wish to bound the effects of the three sources of error separately; therefore it will be convenient to write $|\psi\rangle$ as the sum $|\psi_L\rangle + |\psi_N\rangle + |\psi_D\rangle$, where

$$\begin{aligned} |\psi_L\rangle &= 2^{-pn/2} \sum_{h \in B_G} e^{2\pi i \lambda (f(x) + \mu \nabla f(x) \cdot (h - g_0))} |h\rangle, \\ |\psi_N\rangle &= 2^{-pn/2} \sum_{h \in B_G} \left[e^{2\pi i \lambda f(x + \mu(h - g_0))} - e^{2\pi i \lambda (f(x) + \mu \nabla f(x) \cdot (h - g_0))} \right] |h\rangle \\ &= 2^{-pn/2} \sum_{h \in B_G} \left[e^{2\pi i \lambda f(x + \mu(h - g_0))} - e^{2\pi i \lambda (f(x + \mu(h - g_0)) - \epsilon_N(x, h))} \right] |h\rangle, \\ |\psi_D\rangle &= 2^{-pn/2} \sum_{h \in B_G} \left[e^{2\pi i \lambda c_{r \circ c_f \circ c_p}(c_d(x), h)} - e^{2\pi i \lambda f(x + \mu(h - g_0))} \right] |h\rangle \\ &= 2^{-pn/2} \sum_{h \in B_G} \left[e^{2\pi i \lambda (f(x + \mu(h - g_0)) + \epsilon_D(x, h))} - e^{2\pi i \lambda f(x + \mu(h - g_0))} \right] |h\rangle. \end{aligned}$$

Noting that, for any real α and β ,

$$|e^{2\pi i \lambda (\alpha + \beta)} - e^{2\pi i \lambda \alpha}| = 2|\sin \pi \lambda \beta| \leq 2\pi \lambda |\beta|,$$

we have

$$\begin{aligned} \|\psi_N\|_2 &\leq 2^{-pn/2} \sqrt{\sum_{h \in B_G} (\pi \lambda M \mu^2 (h - g_0) \cdot (h - g_0))^2} \\ &= 2^{-pn/2} \pi \lambda M \mu^2 \sqrt{2^{dn} \left(\frac{16^n}{80} - \frac{4^n}{24} + \frac{2^n}{12} + \frac{7}{240} - \frac{1}{12 \times 2^n} \right)} \\ &\leq 4^{n-1} \pi \lambda M \mu^2 / \sqrt{5} \\ &\leq (1 - \epsilon)/3 \end{aligned}$$

by (4), and

$$\|\psi_D\|_2 \leq 2^{-pn/2} \sqrt{\sum_{h \in B_G} (2\pi\lambda\nu)^2} = 2^{-pn/2} \times 2\pi\lambda\nu\sqrt{|B_G|} = 2\pi\lambda\nu \leq (1 - \epsilon)/3$$

by (5).

Next we consider the error introduced by ‘frequency leakage’. If the components of $\nabla f(x)$ are integer multiples of $1/2^n \lambda\mu$, then $U_{QFT}|\psi_L\rangle$ is equal to a computational basis state, identifying $\nabla f(x)$ exactly. In the general case, we obtain instead a superposition, which strongly weights computational basis states representing gradients close to $\nabla f(x)$. We have

$$\begin{aligned} U_{QFT}|\psi_L\rangle &= 2^{-pn} \sum_{g \in B_G} \sum_{h \in B_G} e^{2\pi i(g \cdot h/2^n + \lambda(f(x) + \mu \nabla f(x) \cdot (h - g_0)))} |j\rangle \\ &= e^{2\pi i \lambda(f(x) - \mu \nabla f(x) \cdot g_0)} \bigotimes_{m=1}^p |\phi_m\rangle, \text{ where} \\ |\phi_m\rangle &= 2^{-n} \sum_{g_m=0}^{2^n-1} \sum_{h_m=0}^{2^n-1} e^{2\pi i h_m \left(\frac{g_m}{2^n} + \lambda\mu \frac{\partial f(x)}{\partial x_m} \right)} |g_m\rangle \\ &= 2^{-n} \sum_{g_m=0}^{2^n-1} \frac{1 - e^{2\pi i \left(g_m + 2^n \lambda\mu \frac{\partial f(x)}{\partial x_m} \right)}}{1 - e^{2\pi i \left(\frac{g_m}{2^n} + \lambda\mu \frac{\partial f(x)}{\partial x_m} \right)}} |g_m\rangle. \end{aligned}$$

The factors $|\phi_m\rangle$ are state vectors of unit magnitude, and note that

$$\begin{aligned} |\langle g_m | \phi_m \rangle| &\leq \frac{2^{1-n}}{\left| e^{\pi i \left(\frac{g_m}{2^n} + \lambda\mu \frac{\partial f(x)}{\partial x_m} \right)} - e^{-\pi i \left(\frac{g_m}{2^n} + \lambda\mu \frac{\partial f(x)}{\partial x_m} \right)} \right|} \\ &= 2^{-n} \left| \csc \left(\pi \left(\frac{g_m}{2^n} + \lambda\mu \frac{\partial f(x)}{\partial x_m} \right) \right) \right| \\ &= 2^{-n} \left| \csc \left(\pi \lambda\mu \left(c_{g,m}(g_m) - \frac{\partial f(x)}{\partial x_m} \right) \right) \right| \quad (9) \\ &\leq 2^{-n} |\csc(\pi \lambda\mu \delta)| \quad (10) \end{aligned}$$

whenever $|c_{g,m}(h_m) - \partial f(x)/\partial x_m| \geq \delta$; we obtain (9) because $\lambda\mu c_{g,m}(g_m) + g_m/2^n$ is always an integer, and $|\csc|$ is even and has period π ; we obtain (10) due to the shape of the cosecant function and because, by (6),

$$\pi \lambda\mu (c_{g,m}(g_m) - \partial f(x)/\partial x_m) \in (-\pi + \pi \lambda\mu \delta, \pi - \pi \lambda\mu \delta).$$

Note that the projection P can be written as $P_1 \otimes \cdots \otimes P_d$, where

$$P_m = \sum_{h_m=0}^{2^n-1} \left\{ |h_m\rangle \langle h_m| : \left| c_{g,m}(h_m) - \frac{\partial f(x)}{\partial x_m} \right| < \delta \right\}.$$

Then

$$\begin{aligned} \|PU_{QFT}|\psi_L\rangle\|_2 &= \left| e^{2\pi i \lambda(f(x) - \mu \nabla f(x) \cdot g_0)} \right| \prod_{m=1}^p \|P_m |\phi_m\rangle\|_2 \\ &= \prod_{m=1}^p \sqrt{1 - \|(I - P_m)|\phi_m\rangle\|_2^2} \end{aligned}$$

$$\begin{aligned}
&= \prod_{m=1}^p \sqrt{1 - \sum_{h_m=0}^{2^n-1} \left\{ |\langle h_m | \phi_m \rangle|^2 : \left| c_{g,m}(h_m) - \frac{\partial f(x)}{\partial x_m} \right| \geq \delta \right\}} \\
&\geq \prod_{m=1}^p \sqrt{1 - 2^n (2^{-n} |\csc(\pi \lambda \mu \delta)|)^2} \\
&= (1 - 2^{-n} \csc^2(\pi \lambda \mu \delta))^{p/2} \\
&\geq (2 + \epsilon)/3,
\end{aligned}$$

by (8).

By the triangle inequality,

$$\|P|\chi\rangle\|_2 \geq \|PU_{QFT}|\psi_L\rangle\|_2 - \|PU_{QFT}|\psi_D\rangle\|_2 - \|PU_{QFT}|\psi_D\rangle\|_2.$$

Since U_{QFT} is an isometry and P is a projection and therefore a contraction,

$$\|P|\chi\rangle\|_2 \geq \frac{2 + \epsilon}{3} - \frac{1 - \epsilon}{3} - \frac{1 - \epsilon}{3} = \epsilon.$$

□

6 Some implementation and efficiency considerations

Theorem 1 established that the algorithm $A(n, \nu, \lambda, \mu; x)$ can perform at any required level of precision, given suitable operating parameters. The algorithm consists of two quantum Fourier transforms, a phase rotation operator, and the two operations U_+ and U_f together with their inverses.

Because we have restricted the sampling grid side-length to powers of two, the quantum Fourier transform is easily computed. The standard quantum Fourier transform in a 2^n -dimensional state space uses just n Hadamard gates and $n(n-1)/2$ controlled phase rotation gates; see [1] for details. The p -dimensional quantum Fourier transform U_{QFT} required by the gradient estimation algorithm is simply the p th tensor power of the standard quantum Fourier transform, meaning that it can be implemented by applying the standard quantum Fourier transform, simultaneously but independently, to the p factors of \mathcal{G} .

The difficulty of implementing the phase rotation operator U_R depends on the data storage method used for function values, i.e., on $(\mathcal{R}, B_{\mathcal{R}}, c_r)$. Implementation is straightforward if a binary fixed-point representation is used, that is, if \mathcal{R} is the state space of a system of say N bits, and the bit sequence

$$(r_{N-1}, \dots, r_1, r_0)$$

represents the value $a_0 + a_1 \sum_{k=0}^{N-1} 2^k r_k$, for some real constants a_0 and a_1 . In this case we can simply pass each bit independently through a phase rotation gate, with matrix representation

$$\begin{pmatrix} 1 & 0 \\ 0 & e^{2\pi i \lambda a_1 2^k} \end{pmatrix};$$

these phase rotation gates are similar to, but simpler than, the controlled phase rotation gates used in the quantum Fourier transform.

The operators U_+ and U_f and their inverses simply perform machine arithmetic. The operators U_+ and U_+^{-1} each involve one multiplication and one addition per domain dimension. The complexity of these operations in gate operations depends on the precision required.

The computational complexity of the operator U_f is entirely dependent on the given function f . It is usual in complexity analyses of computations involving a black-box function f to assume that evaluations of f will be the dominating cost, measuring complexity by counting function evaluations. By that measure, the gradient estimation algorithm scores very well, as it requires two oracle operations, U_f and U_f^{-1} , the latter having presumably the same complexity as the former. (In fact, recall that Section 3 assumed $B_{\mathcal{R}}$ to be a group; if this group is taken to be \mathbb{Z}_2^N , that is, if the computed value is stored in \mathcal{R} using the XOR operation, then U_f^{-1} is just U_f .)

Of course, in order to perform at the required level of precision, we may require very great accuracy in the evaluations of f , and in the other computations. Note that this is a universal feature of machine computation.

7 Conclusion

Theorem 1 shows that the gradient of a real-valued multivariate function can be evaluated to any required accuracy using just two function evaluations. As with any digital computation, increased accuracy in the answer requires increased precision during the computation. Thus the quantum complexity of the gradient estimation problem is constant in dimension, which compares favourably with the classical complexity, which is linear in dimension, for very high-dimensional functions.

Reference

- [1] NIELSEN, M. A., and CHUANG, I. L., *Quantum Computation and Quantum Information*, Cambridge University Press, 2000.